Investigating IPv6 Traffic

What happened at the World IPv6 Day?

Nadi Sarrar[1], Gregor Maier[2], Bernhard Ager[1], Robin Sommer[2,3], and Steve Uhlig[4]

[1] TU Berlin / Deutsche Telekom Laboratories, Berlin, Germany
[2] International Computer Science Institute, Berkeley, CA, USA
[3] Lawrence Berkeley National Laboratory, Berkeley, CA, USA
[4] Queen Mary, University of London, London, UK

Abstract. While the IETF standardized IPv6 more than fifteen years ago, IPv4 is still the prevalent Internet protocol today. On June 8th, 2011, several large content and service providers coordinated a large-scale IPv6 test-run, by enabling support for IPv6 simultaneously: the World IPv6 Day. In this paper, we compare IPv6 activity before, during, and after the event. We examine traffic traces recorded at a large European Internet Exchange Point (IXP) and on the campus of a major US university; analyzing volume, application mix, and the use of tunneling protocols for transporting IPv6 packets.

For the exchange point we find that native IPv6 traffic almost doubled during the World IPv6 Day while changes in tunneled traffic were limited. At the university, IPv6 traffic increased from 3–6 GB/day to over 130 GB/day during the World IPv6 Day, accompanied by a significant shift in the application and HTTP destination mix. Our results also show that a significant number of participants at the World IPv6 Day kept their IPv6 support online even after the test period ended, suggesting that they did not encounter any significant problems.

1 Introduction

The fourth incarnation of the Internet Protocol (IPv4) successfully supported the phenomenal growth of the Internet since its introduction in 1981. Yet, due to this unexpected success, the pressure from the IPv4 address space exhaustion is being felt more and more. This led to the standardization of IPv6 more than 15 years ago, which provides a significantly larger address space. Since then, the transition from IPv4 to IPv6 is happening at a lethargic pace. One of the reasons for the hesitant adoption of IPv6 by end-users is the limited amount of content available through IPv6. A reason for network operators is the fear of breaking critical services. Indeed, the current best practices for deploying IPv6, such as white-listing of well-known network regions, are very conservative. Furthermore, such approaches prevent us from gaining insights into the challenges involved with a global transition to IPv6.

To fill the gap, several operators coordinated a joint experiment on June 8th, 2011: the World IPv6 Day. For the duration of that day, the participants agreed to enable IPv6 support in parts of their networks. Participants included Comcast, Google, Facebook, Microsoft, and many others. Their observations have been reported at the IETF 81 meeting. They found that besides a significant and sustained increase of IPv6 traffic on and

Table 1. Overview of data sets. All data sets are from 2011.

Name	Type	Location	Start date	Duration
JUN1	Packet	Campus	Thu, Jun 2	9 d
JUN2	Packet	Campus	Fri, Jun 17	4 d
JUN3	Packet	Campus	Fri, Jun 24	7 d
IXP1	sFlow	IXP	Wed, Jun 1	22 d
IXP2	sFlow	IXP	Mon, Aug 8	7 d

after the World IPv6 Day, the awareness of IPv6 increased dramatically, and the experience obtained through real IPv6 deployments and measurements were invaluable. The presented results were focused mainly on operational questions, e. g., bandwidth, number of clients, and "IPv6 brokenness".

In this paper, we complement these observations by investigating IPv6 traffic characteristics from two vantage points in the Internet. We examine the use of tunneled IPv6, the presence of applications in IPv6 traffic, and highlight the major IPv6 traffic contributors in the Internet. Our study is based on two traces of production Internet traffic. The first was collected at a large European Internet Exchange Point (IXP) interconnecting hundreds of networks. The second has been gathered at a major US university, a fundamentally different vantage point compared to the IXP, both in scale and level of traffic aggregation. Combined, the two data sets enable us to take a broad look at the impact of the World IPv6 Day.

To the best of our knowledge, this paper is the first systematic study of what has happened around the World IPv6 Day. Our contributions include characterizations of:

Traffic volume: In both traces, we observe a steep and sustained increase of IPv6 traffic. Native IPv6 traffic doubled at the IXP and increased more than 20-fold at the campus.

Tunneling mechanisms: Encapsulated packets contribute a large fraction of IPv6 traffic at the IXP. Teredo tunnels are widespread but mostly idle.

Application mix: Since the World IPv6 Day, the application mix of native and 6in4 IPv6 traffic changed fundamentally and now exhibits similarities to IPv4.

Traffic contributors: Since the World IPv6 Day, YouTube is the main contributor at the campus vantage point. A large content provider is the main contributor at the IXP.

The remainder of this paper is organized as follows. In Section 2, we provide details about our two data sets. We investigate overall IPv6 traffic volume and tunnel encapsulations in Section 3 and the application mix in Section 4. In Section 5 we identify the content providers that contribute most traffic before, during, and after the World IPv6 Day. We present related work in Section 6 and summarize our results in Section 7.

2 Data Sets

We base our analysis on network traffic gathered at the Internet uplink of a major US university and at a large European Internet Exchange Point (IXP). Table 1 gives an overview of our data sets.

In addition to analyzing native IPv6 traffic, we also investigate commonly used tunnel encapsulation methods to transfer IPv6 datagrams over IPv4. In particular, we analyze *Teredo* (RFC 4380), *6in4* (RFC 4212), and *AYIYA*[5] encapsulations. We note that 6in4 encapsulation also covers *6to4* (RFC 3068) and *6rd* (RFC 5969). Some tunneled traffic can be detected by filtering on a specific UDP port; Teredo uses UDP port 3544, and AYIYA commonly runs on port 5072. In contrast, 6in4 has its own IP protocol numer, 41, which can be used for filtering. In all of our analyses, we further verified that the tunnel payload actually contains an IPv6 packet to mitigate against false positives.

Internet Exchange Point: The IXP data sets consist of anonymized sampled sFlow records from the whole traffic exchanged at the IXP. More than 400 networks currently exchange traffic at this IXP. sFlow does not employ flow record aggregation like NetFlow. Instead, sFlow samples one out of n packets and exports the initial portion of it as a sFlow record. The sFlow probes at the IXP use a sampling ratio of 1:2^{14}. We use a customized version of `sflowtool` [14] to extract relevant portions from the sFlow data. As a sFlow record corresponds to the initial portion of a packet, it is possible to examine the protocol and tunneling stack.

US university: We base our analysis of IPv6 traffic at the US university campus on packet level traces collected at the university's central uplink to the Internet. We limited the trace collection to native IPv6 traffic, 6in4 encapsulated traffic and IPv4 traffic on Teredo's well-known UDP port. We then analyze these traces using a customized version of the Bro IDS [13] capable of analyzing tunneled IPv6 traffic.

3 Traffic volume and tunneling

We start by investigating the overall volume of IPv6 traffic before, during, and post the World IPv6 Day. This enables us to calibrate our expectations for subsequent analyses when we dig deeper into used protocols, applications, and traffic contributors.

In Figure 1 and Figure 2 we plot the total bandwidth of IPv6 traffic (native and tunneled) over time at the IXP and the US university, respectively. The World IPv6 Day is highlighted by a gray bar. We observe that before the official start of the World IPv6 Day (at midnight UTC), IPv6 traffic begins to ramp up as content providers enable IPv6 on their systems. During the World IPv6 Day, we observe a 30 % increase of IPv6 traffic at the IXP and an increase from 3–6 GB/day to over 130 GB at the university. We also find that the IPv6 traffic volume remains high after the World IPv6 Day officially ended, indicating that a significant number of participants kept their IPv6 support enabled, and suggesting that they did not encounter significant problems. This is consistent with other reports [1, 9, 15] that observed similar behavior during and after the World IPv6 Day.

Analyzing IPv6 traffic in 1 hour bins shows a clear time-of-day pattern (plot not shown). During and after the World IPv6 Day, the traffic volume during the busy-hour has increased significantly while the traffic dips during off-hours has remained unchanged, indicating that only peak usage has changed but not baseline activity.

We next turn to the question of how much IPv6 traffic is tunneled versus native IPv6 traffic. At the university campus we find hardly any tunneled traffic. At the IXP tunneled traffic is more common. In Figure 3, we plot the IPv6 volume by tunnel encapsulation type for the IXP data sets. During and after the World IPv6 Day, we observe a significant

[5] `http://www.sixxs.net/tools/ayiya/`

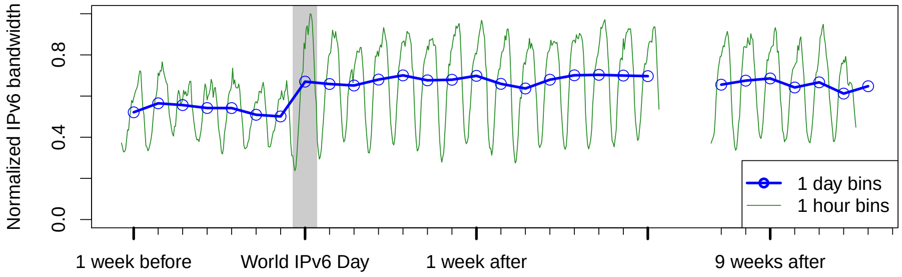

Fig. 1. Total IPv6 traffic volume (IXP). The tick marks are at noon UTC.

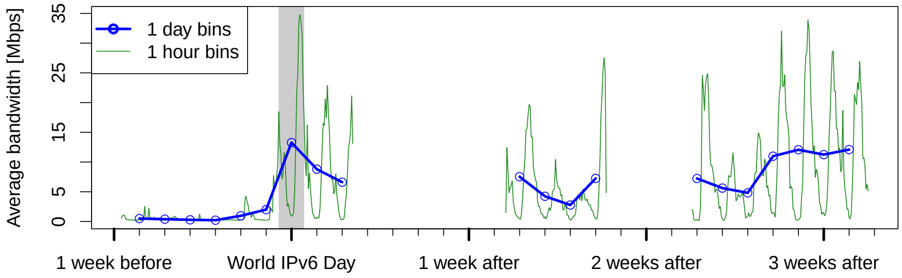

Fig. 2. Total IPv6 traffic volume (campus). The tick marks are at noon UTC.

increase in native IPv6 traffic, while tunneled traffic remains essentially unchanged. The fraction of tunneled traffic decreases accordingly from 69 % to 58 % on average.

We next compare the packet size distributions of IPv4 and IPv6 traffic during the World IPv6 Day and plot the results in Figure 4. We remove the tunnel headers and plot the size of the innermost IPv4 or IPv6 packet. IPv4 shows the usual distribution with peaks at small packet sizes (32 %) and large packet sizes ≥1,492 bytes (25 %). The packet size distribution for IPv6 at the US university resembles the one of IPv4. However, since an IPv6 header is larger than an IPv4 header without options, we find that the "small" packets for IPv6 are slightly larger. We also observe an additional mode at 1,280 bytes for IPv6. This represents the minimum MTU for IPv6 (RFC 2460), and the recommended MTU for tunneling mechanisms in order to mitigate problems with fragmentation (RFC 4380, RFC 4212). We observe a different packet size distribution for IPv6 at the IXP that shows a significantly larger fraction of small packets. More than 82 % of all IPv6 packets are at most 72 bytes in size. Moreover, we notice two modes in the distribution of larger packets, one at the full MTU, and another one at 1,280 bytes. The latter is more pronounced than in the campus data set.

To understand what causes this disparity, we take a closer look at the IPv6 packet size distribution at the IXP by breaking it down according to the type of packet encapsulation. Figure 5 compares the IPv6 packet size distributions for native, 6in4, Teredo, and AYIYA packets. We find strong differences between different encapsulation techniques. Native IPv6 traffic is the only significant source of full-sized 1,500 byte packets, since tunneled traffic needs room for additional encapsulation headers. In contrast to the native IPv6 traffic in the campus data set, we still observe larger fractions of small pack-

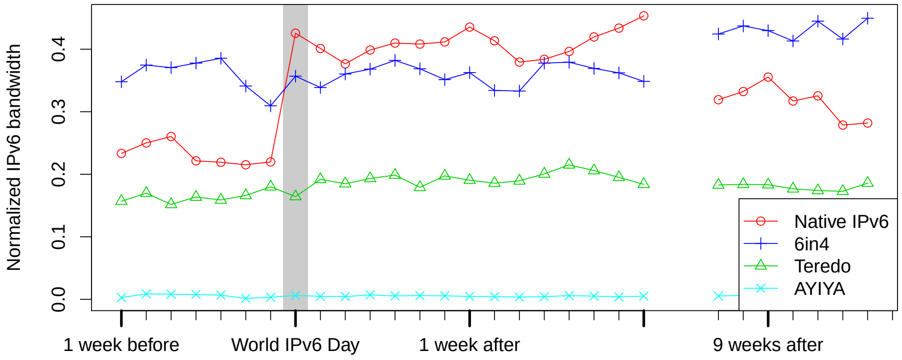

Fig. 3. IPv6 traffic volume by tunnel encapsulation (IXP).

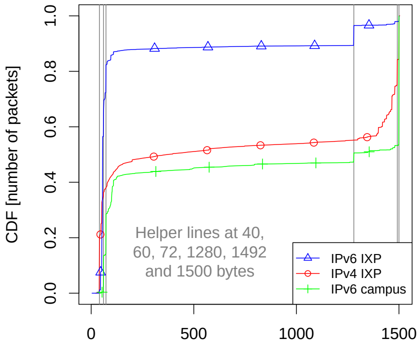

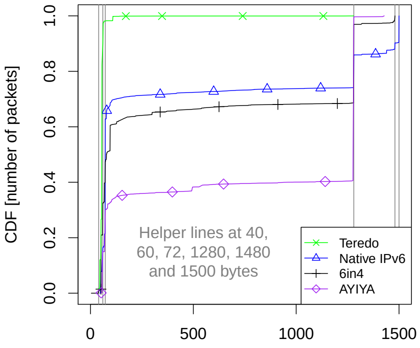

Fig. 4. Packet size distributions of IPv4 and IPv6 traffic (IXP and campus).

Fig. 5. Packet size distributions per encapsulation type (IXP).

ets and a stronger mode at 1,280 bytes. While the packet size distributions for native, 6in4, and AYIYA traffic show some similarities to IPv4, we find that 98 % of Teredo packets are small. A closer examination reveals that at our vantage point, Teredo is mostly composed of control traffic: 76 % of all observed Teredo packets are keep-alive messages (IPv6 headers without payload), and 23 % are ICMP messages. Since Teredo contributes 62 % of IPv6 packets during the IPv6 day, we conclude that Teredo skews the overall packet size distribution dramatically.

4 Application mix

We now turn to the application layer protocol mix of IPv6 traffic. We utilize Bro's dynamic protocol detection framework [4] to classify application layer protocols in the university data sets. As the IXP data set only provided sampled packet headers, we rely on well-known port numbers to identify applications. We use a selection of 86 well-known ports which have been shown to work reasonably well [11]. We report the top

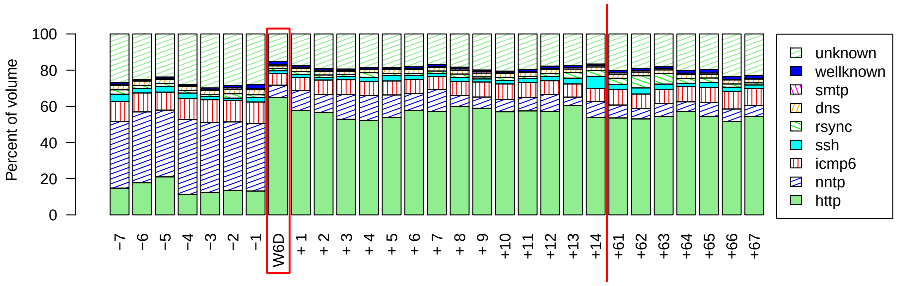

Fig. 6. Application mix per day for *native* IPv6 traffic (IXP).

protocols and aggregate other traffic on well-known ports into the category *well-known*. If the port numbers do not allow to infer the application layer protocol, we attribute the traffic to the *unknown* category.

Figure 6 shows the daily application mix for native IPv6 traffic at the IXP for IXP1 and IXP2. The World IPv6 Day is highlighted by a red rectangle and IXP1 and IXP2 are separated by a red vertical line. Prior to the World IPv6 Day, NNTP was the strongest contributor with about 40 % of the volume, a protocol now frequently used for file-sharing [7]. While we cannot reliably identify P2P traffic in the IXP dataset, its share must be less than 30 % (sum of "well-known" and "unknown" categories). In contrast, Labovitz [8] reports P2P as the main contributor in IPv6 traffic before the World IPv6 Day, with 61 % of the total volume. ICMPv6 contributes 10 % to 13 % of the overall traffic volume. During the World IPv6 Day, the application mix has changed substantially. HTTP is dominating with more than 60 % of the traffic volume, NNTP dropped to 7 % and ICMP to 6 %. In addition, "unknown", and "well-known" now account for less than 15 %. After the World IPv6 Day, the application mix stays roughly similar to the one during the World IPv6 Day, with HTTP loosing about 7 to 10 % of its popularity and ICMPv6 slowly rising up to 9 %.

In Figure 7, we plot the application mix for the campus data sets. We again highlight the World IPv6 Day with a red rectangle and separate different traces with a vertical line. Similar to the IXP we notice a strong shift in the application mix during and after the World IPv6 Day. Before the World IPv6 Day, DNS traffic is in general the main contributor. During the ramp-up to the World IPv6 Day, at and post the World IPv6 Day, we see that HTTP is dominating with a share of up to 97 %. The DNS traffic volume remains unchanged (1–2 GB/day), indicating that it is caused by server-to-server DNS communication and not client requests.

Inside 6in4 tunnels: Since we separately observe multiple different IPv6 tunneling mechanisms at the IXP, we next analyze a breakdown of the application mix according to the tunneling protocol. However, we discuss only 6in4 tunnels since Teredo is almost entirely control traffic and AYIYA lacks volume to provide meaningful results. In 6in4 traffic, which is responsible for more than 32 % of the volume, the most prevalent packets are IPv6 fragments. Further examination of these fragments reveals that half of them have a size of 1,280 bytes (at offset 0), while the other half has 96 bytes. Almost all of the fragments use UDP as transport protocol. We investigated the fragments with offset

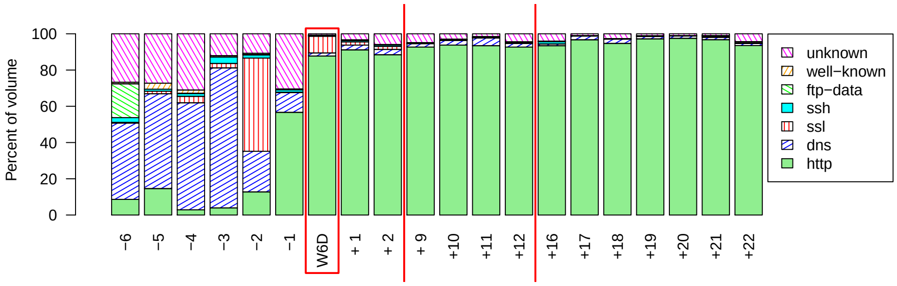

Fig. 7. Application mix per day for *all* IPv6 traffic (campus).

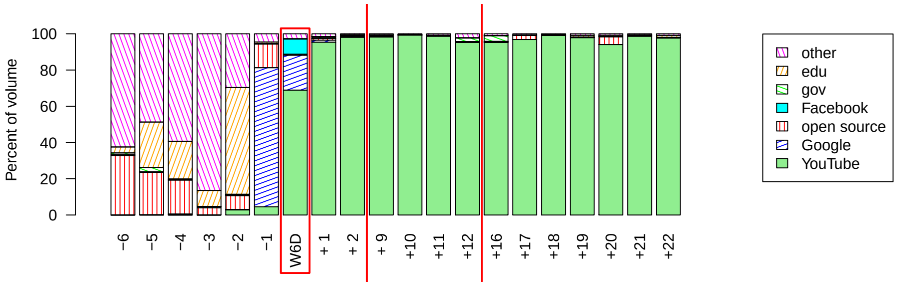

Fig. 8. Daily HTTP mix (campus).

0 to get the UDP port numbers, which appear to be random. Assuming these fragments belong together, the size of the original IPv6 packet before fragmentation would have been 1,320 bytes, which is the minimum IPv6 MTU of 1,280 plus the size of an IPv6 header. We speculate that a broken client software tried to send packets with minimum MTU to prevent fragmentation but forgot to account for the IPv6 header. Before the World IPv6 Day, HTTP was typically at 1–5 % of the traffic volume. During and after the World IPv6 Day, the HTTP fraction increases to 10–16 %. Unknown traffic is at 45 % before and at 52 % during and after the World IPv6 Day.

5 Traffic sources

Since HTTP dominates in the campus environment (up to 97 % of total volume), we analyze HTTP in more detail. We utilize Bro's HTTP analyzer and extract the HTTP server name from the Host header field. We use this information to group HTTP requests by their destination (e.g., YouTube) and plot the result in Figure 8. The "open source" category consists of HTTP-enabled open source software sites, including `freebsd.org`, `mozilla.com`, and `ubuntu.com`. The "gov" and "edu" categories contain all sites under their respective top level domains.

We find that the mix of popular HTTP sites varies from day to day before the World IPv6 Day. Open source and edu sites have significant shares and a large fraction of the traffic is generated by "other" sites. During and after the World IPv6 Day, we observe a significant change with YouTube and Google being responsible for most IPv6 HTTP

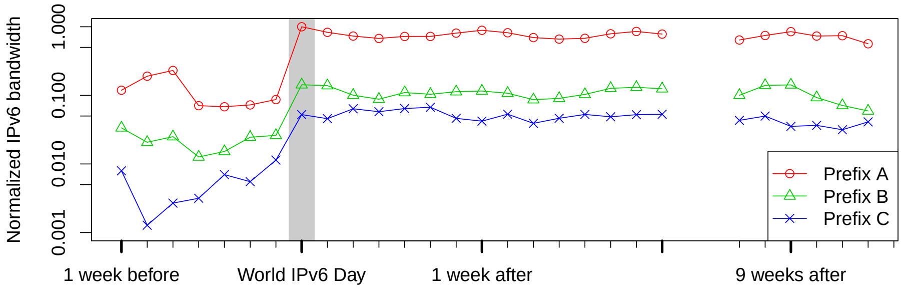

Fig. 9. Normalized bandwidth of *all* IPv6 traffic per prefix before, at, and post the World IPv6 Day (IXP). The plot shows three examples out of the top ten high-volume IPv6 prefixes. Note the log-scale y-axis.

traffic. According to our data, Google enabled IPv6 just before the official start of the World IPv6 Day and disabled IPv6 again after the World IPv6 Day. In contrast, we observe that YouTube kept IPv6 enabled after the World IPv6 Day. Considering that HTTP dominates the application mix and YouTube dominates the HTTP mix after the World IPv6 Day, we conclude that a large volume of IPv6 traffic after the World IPv6 Day is contributed by YouTube.

At the IXP we see more than 3,500 unique IPv6 prefixes. We investigate the largest prefixes in terms of IPv6 traffic volume. Figure 9 shows three out of the top 10 prefixes from the World IPv6 Day. With the help of the IXP we were able to identify prefix A as belonging to a large content provider, and prefixes B and C as large IPv6 enabled stub networks. Only prefix A is actively participating on the World IPv6 Day. Yet, all of them see a roughly ten-fold sustained increase in traffic volume since the IPv6 day. This highlights that passively participating networks can exhibit as much of a change as actively participating ones.

6 Related Work

To the best of our knowledge this paper is the first to perform a systematic study of the IPv6 traffic around the World IPv6 Day. However, there are a number of reports about IPv6 and the World IPv6 Day in the proceedings of the IETF 81 meeting in Canada, July 2011, contributed by the Operations and Management working group.

Palmer and Thaler from Microsoft provide an experience report [12] about the IPv6 activation of several Microsoft's domains. They report having only few connectivity issues. Windows Vista and Windows 7 dominate the observed system types. 91 % of the connections were native IPv6, and less than 1 % were using Teredo. This is consistent with our results about the idleness of Teredo tunnels, and also surprising since Microsoft has enabled Teredo tunneling as a default service since Windows Vista.

Bob Hinden from Check Point reports in [6] about their experience of enabling IPv6 for their company website by using load balancers to handle IPv6. They encountered less difficulties than expected and kept IPv6 active after the World IPv6 Day.

Comcast provides a summary of their IPv6 experiences in [1]. Comcast deployed SMTP over IPv6 by duplicating their infrastructure. Consistent with our results, they report a significant sustained increase of IPv6 traffic at the World IPv6 Day.

In contrast to this study, the above reports were limited to either a few web sites of a single operator, or in case of Comcast to a set of test customers. Still, the reported IPv6 traffic trends and conclusions are consistent with our results.

Hurricane Electric is an early IPv6 adopter—they enabled IPv6 in 2001. Similar to other reports, they observed [9] an IPv6 traffic increase during and after the World IPv6 Day. They also report on path MTU problems and ICMPv6 blocking caused by too aggressive filtering. In addition, they find that 11 % of ASes are present in the IPv6 routing table in August 2011, up from 3.6 % three years earlier.

Wijnen et al. [15] present results from active measurements including DNS, ping6, traceroute6, and HTTP probes. The data was gathered from 40 different vantage points from June 1st through June 11th, 2011. For example, they performed DNS AAAA queries to participating websites and found that nearly all World IPv6 Day participant web sites could be resolved successfully from all of their vantage points. Interestingly, the results also indicate effects due to negative caching of DNS records, as a number of vantage points were not able to resolve AAAA records of some participant, while other vantage points were. Furthermore, they show that after the World IPv6 Day, a number of web sites disabled IPv6 connectivity immediately, while DNS servers continued to return AAAA records for as long as half a day.

Claffy [3] provides an extensive survey of available data that enables tracking of IPv6 deployments, performs comparisons with IPv6 at the topology and the DNS level, and calls out to researchers and industry to provide more data. With our paper, we can contribute to some of the areas identified by Claffy, in particular utilization at access and interconnection links, application mix, and IPv6 tunneling.

Labovitz [8] performs a pre World IPv6 Day study of IPv6 traffic across several providers and presents an application mix including tunneled traffic in which P2P traffic dominates. In addition, our paper characterizes how different tunneling protocols are being used.

Cho et al. [2] performed a very early study of IPv6 path problems and latencies compared to IPv4. Limoncelli et al. [10] compare rollout strategies for IPv6. Guérin et al. [5] model incentives in IPv6 deployment.

7 Summary

In this paper, we conduct the first systematic analysis of IPv6 traffic around the World IPv6 Day. We rely on data collected at two vantage points: a large European Internet Exchange Point and the campus of a major US university. We analyze the traffic volume, application mix, and the use of tunneling protocols for transporting IPv6 packets.

We find that native IPv6 traffic almost doubled during the World IPv6 Day, while changes in tunneled traffic were limited. Teredo tunnels contribute a significant fraction to IPv6 traffic, yet only carry control traffic. We observe significant changes in the application mix during the World IPv6 Day, with the IPv6 application mix becoming similar to the IPv4 one. We find a large amount of fragmented IPv6 packets inside

6in4 tunnels for which broken software is a likely cause. Our results also show that a significant number of participants at the World IPv6 Day kept their IPv6 support online even after the test period ended, suggesting that they did not encounter any significant problems.

8 Acknowledgements

This work was supported in part by a fellowship within the postdoctoral program of the German Academic Exchange Service (DAAD) and by NSF Awards CNS-0855125 and OCI-1032889. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors or originators and do not necessarily reflect the views of DAAD or the National Science Foundation.

References

1. BRZOZOWSKI, J., AND GRIFFITHS, C. Comcast IPv6 Trial/Deployment Experiences, July 2011. Internet-Draft: draft-jjmb-v6ops-comcast-ipv6-experiences-01.
2. CHO, K., LUCKIE, M., AND HUFFAKER, B. Identifying ipv6 network problems in the dual-stack world. In *Proceedings of the ACM SIGCOMM workshop on Network troubleshooting: research, theory and operations practice meet malfunctioning reality* (New York, NY, USA, 2004), NetT '04, ACM, pp. 283–288.
3. CLAFFY, K. Tracking ipv6 evolution: data we have and data we need. *SIGCOMM Comput. Commun. Rev. 41*, 43–48.
4. DREGER, H., FELDMANN, A., MAI, M., PAXSON, V., AND SOMMER, R. Dynamic application-layer protocol analysis for network intrusion detection. In *Proc. USENIX Security Symposium* (2006).
5. GUÉRIN, R., AND HOSANAGAR, K. Fostering ipv6 migration through network quality differentials. *SIGCOMM Comput. Commun. Rev. 40* (June 2010), 17–25.
6. HINDEN, B. Check Point's World IPv6 Day Experience. In *IETF 81 V6OPS*. http://www.ietf.org/proceedings/81/slides/v6ops-2.pptx.
7. KIM, J., SCHNEIDER, F., AGER, B., AND FELDMANN, A. Today's Usenet Usage: Characterizing NNTP Traffic. In *Proc. IEEE Global Internet Symposium* (2010).
8. LABOVITZ, C. Six month, six providers and IPv6. Tech. rep., March 2011. http://www.monkey.org/~labovit/papers/v6sixmonths.pdf.
9. LEVY, M. IETF 81 - World IPv6 Day Operators Review. In *IETF 81 V6OPS*. http://www.ietf.org/proceedings/81/slides/v6ops-19.pdf.
10. LIMONCELLI, T. A., AND CERF, V. G. Successful strategies for ipv6 rollouts. really. *Commun. ACM 54* (April 2011), 44–48.
11. MAIER, G., FELDMANN, A., PAXSON, V., AND ALLMAN, M. On dominant characteristics of residential broadband internet traffic. In *Proc. Internet Measurement Conf. (IMC)* (2009).
12. PALMER, C., AND THALER, D. World IPv6 Day at Microsoft. In *IETF 81 V6OPS*. http://www.ietf.org/proceedings/81/slides/v6ops-1.pptx.
13. PAXSON, V. Bro: A system for detecting network intruders in real-time. *Computer Networks Journal 31*, 23–24 (Dec. 1999), 2435–2463. Bro homepage: http://www.bro-ids.org.
14. InMon: sFlow Toolkit. http://www.inmon.com/technology/sflowTools.php.
15. WIJNEN, B., ABEN, E., WILHELM, R., AND KISTELEKI, R. World IPv6 Day—What did we learn? In *IETF 81 V6OPS*. http://www.ietf.org/proceedings/81/slides/v6ops-4.pdf.